

**A METHOD AND SYSTEM FOR CONFIGURING PROCESSING ENTITIES
ACCORDING TO THEIR ROLES IN A DATA PROCESSING SYSTEM WITH A
DISTRIBUTED ARCHITECTURE**

Technical field

5 The present invention relates to the data processing field, and more specifically to a method and a corresponding system for configuring processing entities according to their roles in a data processing system with a distributed architecture.

Background art

10 Managing configuration changes in a data processing system with a distributed architecture is a time consuming activity, particularly when the system includes a high number of processing entities on which new configurations must be enforced. A typical example is that of a large network with thousands of computers,
15 where software products are periodically upgraded in order to be abreast of the information technology development.

 Software distribution applications have been proposed in the last years to assist a system administrator in efficiently managing deployment of software products from a central site of
20 the network. A software distribution application controls building of packages, which include instructions specifying the actions to be carried out on target computers for installing or removing selected software products; each package further embeds a copy of the software products to be installed. The package is
25 transmitted to the computers, and the corresponding instructions are interpreted so as to enforce the desired software configuration.

However, this solution is not completely satisfactory. Particularly, the administrator is faced with the burden of manually defining the packages to be distributed to the computers. As a consequence, the process of managing the
5 configuration changes is complex and error-prone, particularly if the network includes a high number of computers.

The drawbacks of the solutions known in the art are more acute when the configuration process involves a software application that is distributed across the network. Typically,
10 the distributed application includes different software components for corresponding computers of the network. As a consequence, the task of installing and upgrading the distributed application is very difficult (because of the need to tackle the problem as a whole, without the possibility of addressing a
15 single computer at the time). A typical example is that of a software infrastructure used to manage several aspects of the life cycle of the network itself.

The above-described problems are exacerbated when multiple products are shipped as a single suite; moreover, additional
20 difficulties arise when the network has a complex topology.

Summary of the invention

It is an object of the present invention to provide a method of configuring the entities of the data processing system according to their roles played in the system.

25 It is another object of the present invention to simplify the configuration of the system.

It is yet another object of the present invention to provide a mechanism that is well suited to install and upgrade distributed applications, particularly when the applications
30 include multiple products.

Moreover, it is an object of the present invention to support the configuration of data processing systems with a complex topology.

The accomplishment of these and other related objects is achieved, in a data processing system with a distributed architecture including a plurality of processing entities, each entity playing at least one of a plurality of predetermined roles in the system, by a method of configuring the entities including the steps of: defining a target configuration for each role, identifying the at least one role of each entity, and configuring each entity according to the target configuration corresponding to the at least one role of the entity.

The present invention also provides different computer program applications for performing the method, and respective products storing the program applications.

Moreover, a corresponding data processing system and a server computer for use in the system are encompassed.

The novel features believed to be characteristic of this invention are set forth in the appended claims. The invention itself, however, as well as these and other related objects and advantages thereof, will be best understood by reference to the following detailed description to be read in conjunction with the accompanying drawings.

Brief description of the drawings

- Figure 1a** is a schematic block diagram of a data processing system in which the method of the invention can be used;
- Figure 1b** shows the functional blocks of a generic computer of the system;
- Figure 2** depicts the main software components used for implementing the method;
- Figures 3a-3b** is an activity diagram describing the flow of the method.

Detailed description of the preferred embodiment

With reference in particular to Figure 1a, a data processing system 100 with a distributed architecture is shown. The system 100 consists of a network of computers; the computers are interconnected through an infrastructure 110 (for example, based on a Local Area Network, or LAN), according to a pattern defining a topology of the network 100.

Each computer plays a specific physical role in the network 100; the physical role of the computer depends on the architecture of the network 100. For example, in a Tivoli Management Environment (TME), the network 100 implements an independent region (referred to as Tivoli Management Region, or TMR). The region has a three-tier structure. A TMR server 120s manages the whole region from a central site. The server 120s communicates with several gateways, or Managed Nodes (MN), 120g. The gateways 120g bridge between the server 120s and corresponding clusters of endpoints 120e. Each endpoint 120e defines a target of management actions enforced by the server 120s.

Moreover, each computer (generically denoted with 120) may also play different logic roles specific for applications running

in the network 100. For example, in a software distribution process a computer is configured to be a Source Host (SH), which operates as a preparation and testing central site for software products to be deployed throughout the network 100.

5 A shown in Figure 1b, a generic computer 120 (server, gateway or endpoint) is formed by several units that are connected in parallel to a communication bus 150. In detail, a microprocessor (μ P) 155 controls operation of the computer 120, a Read Only Memory (ROM) 160 stores basic code for a bootstrap of the
10 computer 120, and a Random Access Memory (RAM) 165 is directly used as a working memory by the microprocessor 155. Several peripheral units are further connected to the bus 150 (by means of respective interfaces). Particularly, a mass memory consists of a magnetic hard-disk 170 and a unit 175 for reading CD-ROMs
15 180. Moreover, the computer 120 includes input devices 185 (for example, a keyboard and a mouse), and output devices 190 (for example, a monitor and a printer). A network Interface Card (NIC) 195 is used to connect the computer in the network.

However, the concepts of the present invention are also
20 applicable when the network has a different topology, when the computers are interconnected in another way (for example, through the Internet), or when the network is replaced with an equivalent data processing system with a distributed architecture. Similar considerations apply if the system includes different logic
25 and/or physic entities (such laptops, Personal Digital Assistants (PDAs), mobile telephones, or virtual machines), if the computers have a different structure or include other units. Alternatively, different physic and/or logic roles are envisaged; for example, a computer may operate as a firewall, as a collector of results in
30 an inventory application, and so on.

Considering now Figure 2, a partial content of the working memory of the server in operation is shown. The information (programs and data) is typically stored on the hard-disk and

loaded (at least partially) into the working memory when the programs are running. The programs are initially installed onto the hard-disk from CD-ROM.

An operating system 205 provides a software platform for the server, on top of which other programs can run. A Tivoli Management Framework (TMF) 210 is installed over the operating system 205; the framework 210 defines a software infrastructure that is common to the whole region.

A catalogue 215 specifies a current software configuration of all the computers of the network. For each computer, the configuration catalogue 215 stores a name of its physic role and an inventory of the different software components of the products installed on the computer. The physic role of the computer is set at installation time; the software inventory is updated whenever the software components are installed and/or upgraded on the computer.

The configuration catalogue 215 is accessed by a discovery module 220. The discovery module 220 detects a current level of the framework installed on every computer (exploiting corresponding network services). The discovery module 220 generates a further catalogue 225 dynamically. For each computer, the dynamic catalogue 225 stores the level of the framework and the name of its logic role. The logic role is identified according to the software components installed on the computer; for example, the presence of a module SPEDITOR makes it possible to classify the computer as a source host in a software distribution application.

A driver 230 controls the unit for reading CD-ROMs. The driver 230 is used to upload information relating to software products to be installed in the network (provided on one or more CD-ROMs). For example, the process of the invention is used to install a Tivoli Deployment Suite consisting of a Change Configuration Management (CCM) product, an Application

Performance Management (APM) product, a Tivoli Resource Manager (TRM) product, a Software Distribution (SWD) product, and an Inventory (INV) product.

Each product is provided with a set of images of the 5 different components to be installed on corresponding categories of computers (defined according to their physic and/or logic roles); the images uploaded from the CD-ROMs are stored in a repository 235. Moreover, the product is provided with a file 240 (for example, in the XML format), which describes how to use the 10 different images of the product. The descriptor 240 includes a first section specifying a reference model for the product; for each (physic and logic) role, the reference model includes the name of the component(s) that must be installed on the computers playing that role. Moreover, the descriptor 240 includes another 15 section specifying a transition table for installing the product; for each pair current state/target state of a generic component of the product, the transition table stores an indication of one or more actions required to reach the target state from the current state; an error code is instead specified when the 20 transition from the current state to the target state is not supported.

For example, the transition table specifies that the target state of the component requires its fresh installation (when no previous version of the component is available) or an upgrade 25 (when a former, upgradeable version of the component is already installed); conversely, an error is specified when a direct migration from a very old level of the component to a latest level thereof is not possible. More specifically, let us consider a product including the components COMP1_v.4.1, COMP2_v.3.0, 30 COMP3_v.2.5, and COMP4_v.2.3. The reference model specifying the target states for the generic roles ROLE1, ROLE2 and ROLE3 is:

Role	Targ t state
ROLE1	COMP1_v.4.1
ROLE2	COMP3_v.2.5, COMP4_v.2.3

ROLE3	COMP2_v.3.0
-------	-------------

The descriptor of the product also includes the following transition table:

Target state	Current state	Actions
COMP1_v.4.1	N/A	Install COMP1_v.4.1
	COMP1_v.4.0	Upgrade to COMP1_v.4.1
	COMP1_v.1-COMP1_v.3	Error
COMP2_v.3.0	N/A	Install COMP2_v.3.0
	COMP2_v.2	Upgrade to COMP1_v.3.0
	COMP2_v.1	Error
COMP3_v.2.5	N/A	Install COMP3_v.2.5
	COMP3_v.2.0-COMP3_v.2.4	Upgrade to COMP3_v.2.5
	COMP3_v.1	Error
COMP4_v.2.3	N/A	Install COMP4_v.2.3
	COMP4_v.2.0-COMP4_v.2.2.	Upgrade to COMP4_v.2.3
	COMP4_v.1	Error

5 For example, if the product must be installed on a computer playing the role ROLE3 and wherein the component COMP2_v.2 is already installed, the action "Upgrade to COMP2_v.3.0" must be executed to reach the desired target state.

The configuration catalogue 215, the dynamic catalogue 225,
10 and the descriptor 240 are input to a differencing engine 245. The engine 245 generates a list of actions 250 for enforcing a desired target configuration in the network (as defined in the reference model of the descriptor 240). For each computer, the action list 250 stores one or more records specifying the names
15 of the components to be installed and the actions to be executed for installing the components starting from their current states on the computer. The actions are ordered into a correct sequence; for example, the sequence ensures the inter-operability in the network (installing the different products in a top-down order
20 beginning from the server) and the required dependencies among the products.

The action list 250 is supplied to a plan generator 255. The plan generator 255 exploits the framework 210 for executing some of the actions synchronously. The other actions are chained into one or more plans 260, each one scheduled at a desired time.

5 A module 265 controls the execution of the plans 260 (exploiting the framework 210). For this purpose, the plan executor 265 transmits the list of actions (extracted from each plan 260) and the images of the corresponding components (imported from the repository 235) to the relevant computers.

10 However, the concepts of the present invention are also applicable when the whole application for configuring the network has another structure, when different products must be installed on the computers, or when the current configuration of the computers is detected in another way (for example, only during
15 the configuration process or only dynamically). Similar considerations apply if the products with the corresponding descriptors are provided in a different manner (for example, they are downloaded from a web site of the publisher), if the descriptor has another format, if the actions required to enforce
20 the desired target configuration in the network are executed in a different way, and the like. For example, every computer retrieves the images of the components to be installed from depots distributed across the network; alternatively, the computer accesses the web site of the publisher and then
25 downloads the images directly.

As shown in the activity diagram of Figures 3a-3b, a process 300 of configuring the network described above begins at the black start circle 303 in the swim-lane of the publisher. Continuing to block 306, a new suite of products is made
30 available. At the same time, the publisher defines the corresponding reference models at block 309, and the corresponding transition tables at block 312. The suite of

products with their descriptors is then delivered to the server at block 315.

Considering now block 318 in the swim-lane of the server, the images of the products and the corresponding descriptors are
5 uploaded onto the hard-disk. The installation of the products is typically guided by a wizard developed in the Install Shield Multi-Platform (ISMP) environment. Particularly, at block 321 a panel prompts a system administrator to select the desired products to be installed (with the framework and the plan
10 executor that are mandatory); for every product, the target state for each role (specifying the components to be installed) may be displayed. If the proposed choices are not accepted at block 324, the target configuration of the network may be updated at block 327; the process then returns to block 321.

15 Conversely, as soon as the system administrator confirms the target configuration the process descends into block 330; the physic role of every computer is then identified (from the configuration catalogue). Continuing to block 333, the current configuration of every computer is detected (retrieving the
20 information relating to the level of the respective frameworks dynamically and the information relating to the other components installed on the computer from the configuration catalogue). The computers are then classified at block 336 in the logic role corresponding to the components installed thereon.

25 As soon as the above-described discovery process has been completed, the flow of activity descends into block 339; for each (physic and/or logic) role of every computer, the desired target state is extracted from the reference model; the entry of the transition table associated with the pair current state/target
30 state is then identified. If a migration is not allowed from the current state to the target state, an error condition is entered at block 342. Conversely, the list of actions required to reach

the target state from the current state on the computer is generated at block 345.

Some of the actions resulting from this differencing process are executed synchronously at block 348 under the control of the plan generator (for example, a first installation of the framework or the configuration application running on the server is performed directly). The process then passes to block 351, wherein one or more plans for enforcing the other actions are created and scheduled for execution.

As soon as the scheduled time of a generic plan is reached, the execution of the plan is started at block 354. Continuing to block 357, the list of actions and the corresponding images are transmitted to the relevant computers of the network. In response thereto, an agent running on a generic computer installs the desired components at block 360. Passing to block 363, the result of the installation operation is returned to the server. The server at block 366 updates the record associated with the computer in the configuration catalogue accordingly. The process then ends at the concentric white/black stop circles 369.

However, the concepts of the present invention are also applicable when an equivalent method is performed, or when the process is aborted if a condition preventing its execution is detected (for example, because the process is launched on a computer that is not the server). Similar considerations apply if the space on the hard-disk of the server is checked before loading the images of the products to be installed, or if a mechanism is provided for restarting failed plans (skipping the actions already executed).

The same method is also applicable when a product must be upgraded (for example, by installing a patch). In this case, the patch is shipped with a corresponding new descriptor (including the reference model and the transition table); the descriptor specifies where the patch needs to be applied and how to do so. As described

above, the process detects the current configuration of the computers, and proposes the actions to be carried out for aligning the network to the desired target configuration. As a further enhancement, the server automatically monitors the web site of the publisher; as soon as a new patch is made available, its application is enforced as described above.

More generally, the present invention proposes a method for use in a data processing system with a distributed architecture. The system includes a plurality of processing entities; each entity plays one or more predetermined roles in the system. The method of the invention is designed to configure the entities. The method starts with the step of defining a target configuration for each role. The roles of each entity are identified. Each entity is then configured according to the target configuration corresponding to the roles of the entity.

The solution of the invention makes it possible to configure the entities of the data processing system according to their roles.

As a consequence, the management of any configuration change in the system is strongly simplified.

The proposed mechanism is well suited to cover either the installation or the upgrading of distributed applications, and particularly applications including multiple products. For example, this method is suitable to address a scenario wherein the multi-product application is to be installed in a network wherein some products (possibly of a former version) have already been installed individually; advantageously, the solution of the invention is used to install a software infrastructure used to manage the life cycle of the whole system (even if different applications are contemplated and within the scope of the invention).

Moreover, the proposed method makes it possible to configure data processing systems with a complex topology in a very simple manner.

The preferred embodiment of the invention described above
5 offers further advantages.

Particularly, the roles played by the computers include one or more physic roles.

This characteristic allows the definition of reference models, which are bound by an architecture of the network.

10 Advantageously, the physic role of each computer is retrieved from a memory structure set at the installation of the computer (either stored on the server or on the computer itself).

The proposed mechanism is very simple, but at the same time effective.

15 As a further enhancement, the network also supports one or more logic roles.

This additional feature strongly simplifies the configuration of applications requiring computers with different functions.

20 Preferably, the logic role of each computer is established according to the components (or other equivalent software features) installed on the same.

The devised mechanism is very flexible, and it can be adapted dynamically to the different application requirements.

25 Alternatively, the method only supports physic roles, the physic role of every computer is identified in a different manner (for example, inspecting a directory specifying a topology of the network), the method only supports logic roles, the logic role of every computer is identified in another way (for example,
30 retrieving the information from a register that is set at the installation of the application on the computer), or different roles are envisaged.

In an advantageous embodiment of the invention, a transition table is used to identify the actions required to reach each desired target state from each current state.

This feature makes it possible to automate the whole
5 configuration process.

Preferably, the proposed method is used to install software products (each one provided with a respective reference model and a respective transition table).

The devised mechanism strongly simplifies the maintenance
10 of the software products used in the network.

However, the method of the invention leads itself to be implemented even configuring each computer according to the desired target state in a different manner (for example, without detecting its current state by the server). Alternatively, the
15 proposed method is used to manage different configuration features (such as the content of internal registers or the triggering of monitoring activities).

Advantageously, the solution according to the present invention is implemented with a computer program application,
20 which is provided as a corresponding product stored on a suitable medium.

In a preferred embodiment, the method is performed under the control of the server.

This architecture makes it possible to manage the whole
25 configuration process centrally from a single site.

Alternatively, the program application is pre-loaded onto the hard-disk, is sent to the server through the Internet, is broadcast, or more generally is provided in any other form directly loadable into the working memory of the server. However,
30 the method according to the present invention leads itself to be carried out with an application having a different architecture (for example, implementing an adaptive model), or even with a

hardware structure (for example, integrated in a chip of semiconductor material).

Naturally, in order to satisfy local and specific requirements, a person skilled in the art may apply to the
5 solution described above many modifications and alterations all of which, however, are included within the scope of protection of the invention as defined by the following claims.